CHARTS AND MEASURES [1]

Edited by Andrew McNutt

alt.vis has gone mainstream. The alt.vis workshop was formed as a bastion of alternative thinking and conceptual unsettling, to push the boundaries of visualization, both in research and in practice. By some measures, this has been successful, yielding papers on topics as varied as metal logos, critical practice, sustainability, and so on. Untold truths were told, provocations were placed, dances danced, and so on.

Yet we now arrive at a crisis, alt.vis has now become too mainstream to maintain its alternative purposes. For what else can a venue be thought of that has reached its fourth edition besides that it has become institutionalized? Enured to the real truths and harsh realities of the world surrounding us, alt.vis no longer has the aesthetic clarity or purity to conduct its much-needed work.

To rectify this issue, we propose a new workshop housed within the alt.vis workshop entitled "Charts and Measures". Its intent is to carry alt.vis's heavy burden further and offer a truly alternative venue for visualization research—housed within the now mainstream alt.vis. This new venue will be unbound by the demands and constraints inherent to a workshop growing towards the center. The name of this new venue follows that most well-regarded source of non-satirical information "Shouts and Murmurs" (as featured in venerable scholastic venue the New Yorker). This new venue allows for rich and expressive work as described through extended abstracts, not limited by the bindings of buttoned-down alt venues. We join a long line of previous venues that truly embrace the spirit of alternative work (e.g., SIGBOVIK and Journal of Universal Rejection), which offer true boundary pushing means through which to interrogate the boundaries of alterity.

To ensure the alternative integrity of this new venue, we limit submissions to only those approved to submit by the organizing committee, both of which are composed of Andrew McNutt. A program committee was formed, consisting of relevant experts in the area (Andrew McNutt), and conflicts declared. All papers were deemed as being in conflict, and so the issue was viewed as resolved. We had 36 papers "submitted" from which we selected the 9 that appear here. These proceedings include a range of topics, such as the inherent imperialism of some research fields, to quagmires we have found ourselves struggling through.

DISCOVERY OF COMPUTER SCIENCE BY VISUALIZATION

Big data plays an ever more important role in daily life. The rigorous development and systemization of actionable conceptual models is a growing practice in the visualization field. For instance, there has been a growing systemization of visualization grammars [9] that has led to a variety of successes, such as the prominent Vega-Lite. To this end, we have developed an entirely new discipline in which we apply the methods of traditional science to computers. Through this systematic exploration of this topic, we have developed novel ideas that offer vast potential benefits to visualization.

The first of which is the notion of small machines that enact a series of predefined commands. By combining these predefined commands with a notion of memory held on a long tape (analogous to a blockchain), we believe that many potential useful applications will be unlocked. Moreover, we suggest that these "vis bots" could be useful for a wide variety of domains, including AI-explainability or visualization recommendation. We introduce a number of small guidelines on how they might be used, such as proof of when they will and won't stop—although we leave exploration of the technical details to others. Next is a rigorously defined notion of "computational complexity" which allows for the description and reasoning about the algorithmic performance of these vis bots, which will be vital. Additional work, that, crucially, we will not do, could usefully introduce notion systemizations of these vis bots that could allow practitioners from other fields to use them.

It is possible some of these areas have been touched on by prior research; however, to the best of our knowledge (although we did not check outside of TVCG), we are the discoverers of these likely world-changing topics.

VISUALIZATION HAS A MOUSE PROBLEM

Listen, we've all done it. We've all found ourselves in times of dire straits and found ourselves falling into that easy trap that no one really ought to fall into. We shouldn't be ashamed, but we, as a community, really should stop rebuilding Tableau.

In biological sciences, it is common to focus on one of several model organisms [13] from across the range of kingdoms of life. For instance, Drosophila melanogaster, or fruit fly, is deeply common for its relatively simple genome. Similarly, Mus musculus, or the common mouse, is used throughout fields like immunology to host a wide variety of different experiments. While these models are useful, in that they provide a common medium on which to compare varied experiments, they can lead to over-fixation on a particular model and decline of consideration of others [4].

In visualization, we have picked two "model organisms" that I can see readily: scatterplots (sometimes [7] called a petri dish for perception experiments) and the visual analytics stylings of Tableau. This work forgoes consideration of scatterplots and focuses on the latter of these.

While Tableau works well enough once you are trained in it, I find that practitioners I talk to struggle with the conceptual model presented in Tableau, noting that it impedes just getting the result that they want. This may be because Tableau's (and similar tools) pills-and-shelves purposely centers data exploration, which, as people like Alberto Cairo are prone to pointing out, is not really the task that a lot of people go to visualization for.

While I'm proud of my recreation of Tableau (Ivy [10]) using Tableau as our north star for what makes an effective visualization tool, it limits our imagination (and impact) as a research community. It is, unarguably, a beautiful set of abstractions and is particularly well aligned with one of our few theories, the grammar of graphics. Yet, the total of this as a tool form is complicated and disjoint with many real-world applications—leading to yet another divergence between research and practice.

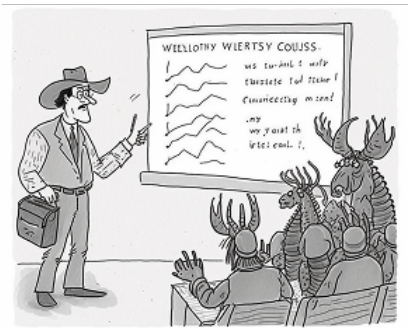

Figure 1: Hi, I'd like to add you to my professional network on LinkedIn [6]

VISUALIZATION FOR VILLAINY 2: THIS TIME I'M MAD ABOUT SOMETHING SPECIFIC

Why would a second paper [11] be necessary to underscore the points that the primary harms we can do with visualization (both as researchers and practitioners) are to be unreflective about our data, tools, and processes? In this work, I will extend our prior work by providing a more focused view on a specific type of evil—rather than focusing on the various types of evil that can be enacted by basically doing nothing and continuing to pursue our basic trajectory as we have been. In the prior work, we enumerated the various ways in which evil might be done, including through indirect means (such as by poisoning our trust in particular chart forms) as well as direct means. This work does not merely reiterate these points and connect them with the resplendent collections of evil (visual and not) that are now omnipresent in the wild, but instead make new points about new things and highlight a specific form of chart-based evil.

I DO NOT CARE ABOUT YOUR SELF-CONGRATULATORY BAR CHART

There is a tendency in visualization (and other fields) for the evaluation to be merely stapled on, in which the system (for most often this happens in systems papers) is evaluated through some rinky-dink study which sports poor experimental methods, unreproducible results, and conclusions which are generally overly optimistic about the value the content. This is then often summarized in a bar chart that highlights how the system's performance is substantially better than everything that came before.

If good papers are about ideas, why do we-gate keep those ideas behind checkbox-adherence to the completion of experimental validations of those ideas? For instance, Voyager's [14] value is not in its user study, but in the empowering way it combines a reified notion of visualization (via Vega-Lite) to enable a recommender. Similar questions might be raised about our use of related work sections, but others [5] have argued that more fiercely.

But how are we (readers) supposed to know if something is good? Personally, I would rather read a thoughtful reflection that extends and complicates the ideas presented. Think about your system from a feminist lens! Consider the cognitive dimensions of your design! Please! Do literally anything other than just show me a bar chart that says some people liked your thing, or that you made a made-up number go up. Failing that: maybe let the literature work it out. If an idea reveals to others how to build on new things (broadly defined), then they will.

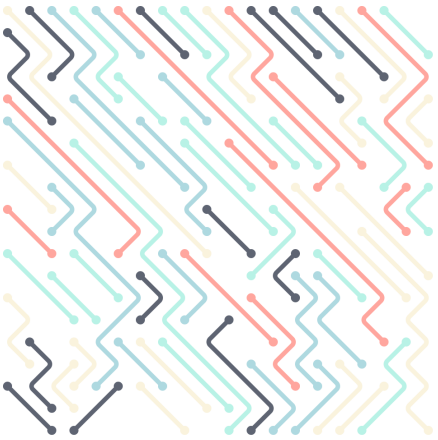

VIS EDUCATION'S EMBRACE OF D3 IS AN ATTEMPT AT SELF JUSTIFICATION

There is sometimes some debate about whether or not visualization is well placed in computer science. Anecdotally: I've had students write in my course reviews "this wasn't really a CS course this was more of a psych course", and I've heard other instructors be questioned about the amount of "computer science" in an upcoming visualization course. While these debates are largely useless and most feel like a way to demean our field and to brand us as non-technical work, it is worth taking a moment to consider how we've internalized these perspectives.

Some parts of this internalized self-loathing can be seen in the presence of terms like "technical HCI" which serve only to create worthless divisions in a field inherently focused on the technical topic of computers. Visualization may have a technical graphics origin [12], but it, subjectively, has spent much of its life span trying to justify its existence (recall Shneiderman's understanding not pictures riff)—leading to something that might be called an inferiority complex—possibly explaining why visualization is so conservative in what it counts as visualization [3].

More concretely: D3 is an extremely common part of visualization courses. I know I have taught it, others have taught it, there are literal textbooks on how to use it, papers and talks describing materials and techniques for making it simpler to use and to teach regularly appear at visualization conferences. There are many extremely powerful things people can do with d3—just go look at the d3 example page [2].

I question how good this actual choice is for helping students understand the ideas we are trying to convey. A visualization course might have learning objectives like "can reason about the meaning of visualization encodings" and "can critique visualizations" or "can construct informative graphics". These are good and useful goals! Yet, the particular collection of idiosyncrasies inherent to d3 provides a massive stumbling block for many trying to approach these tasks. Why do people need to understand the enter-update-exit life cycle's quirks when what they are trying to do is to put a mark there? Sure, there is something to be said for building tools with high ceilings, but usually the dual to that metaphor is low floors. In contrast, courses are sometimes taught with tools (like p5.js) that support artistically intuitive tool metaphors or that use contemporary web tools (like Svelte).

While learning to think in new ways is a valuable part of education, I question whether thinking in the specific framing of d3's vision of functional programming is a useful form of thinking to impart to students. It transfers poorly to other visualization libraries, is out of step with contemporary web tools, and has extremely difficult notions of state to comprehend. Besides building a portfolio piece for students, what benefit does it bring them? If painting courses required bespoke vis bots instead of paint brushes, would we see that choice as appropriate when most painters use paint brushes?

In this essay, I suggest that we are aware of these issues and still keep using d3 as a primary teaching aid, not because it is good, but because it gives visualization a veneer of computer science that we might otherwise lack. d3 is technical, and so teaching it means we're technical, right? Credibility through complication.

LATE BREAKING RESEARCH! VISUALIZATION DISCOVERS AMERICA

Recent advances in geospatial mapping technologies have led to numerous prominent discoveries about land thought to be previously understood. Following these advances, we describe a recently discovered land mass of enormous size, approximately three thousand miles west of France, which we refer to as Vislandia. Initial surveys suggest that this large landmass is approximately 3.797 million square miles. While it is possible that other disciplines might have been aware of this land mass, we believe that it is unlikely that we'd be so far behind cartography yet again (how many times can that possibly happen?). This large land mass offers ample opportunity to develop and consider new areas of visualization research. We suggest that the natural segmentations of the VIS community might be allocated to different sections of this vast zone. By allocating these resources amongst ourselves (ala the scramble for Africa), we might develop an improved area model, by conquering the areas therein.

EXPLORING THE VISUALIZATION LITERACY OF DICE

There has been significant recent excitement towards the exploration of understanding the literacy of large language models. Considerations have been applied towards understanding the relative abilities of different models to complete a varied set of literacy tasks—to a significantly varied set of findings. While these results are exciting, they are deeply entwined with the heat of the AI summer. How are we to look upon this raft of papers in 10 years when technology has changed beyond recognition? Instead, we suggest that investigation of more steadfast technologies will likely be more productive.

In particular, we investigate the visualization literacy of the millennia-old technology of dice. Dice come in many cardinalities (threes, fours, sixes, twenties, and so on), forms (such as black and white and see-through blue), materials (e.g., plastic, stone, wood), and a constellation of usages (conducting games of chance, tricking people through weighted dice, and so on).

We explore how these age-old input devices are able to understand and evaluate visualization using contemporary literacy measures (such as the many VLAT variations). After significantly adapting these tests so that they speak the language of dice, we find that while dice tend to follow a random distribution across the tests, we are strongly heartened that a particular set of dice happened to exhibit extremely high literacy. We suggest that future studies might investigate the "explainability" of these dice to better understand their performance.

HOW MIGHT VIS DIE?

Sometime in the first days of my Ph.D. at an orientation week party, a more senior student asked what field I was in. I said visualization, and he immediately had lots of questions about it. One that we got stuck on was "where does the field expect to be in 10 years?" He really would not accept the answer "that's not how HCI fields work" and eventually settled himself with the idea that the project of visualization research was to visualize ever higher-dimensional data. It is now nearly a decade later, and after most of a decade of ennui, I still struggle with a somewhat basic question:

What is the point of visualization (as an academic field)?

For sure, there have been years of grand challenges papers—although it is questionable how much progress we've been making, as tasks like "visualize uncertainty" and "embrace practitioners" have been appearing on such lists for nearly 20 years [8]—and a variety of success stories—I often refer to 2019 as "year of the BI tool" as Tableau and Looker both got bought for billions of dollars.

But these are how we live: it is the work we do and strive for. While some would say that what matters in life is the journey, not the destination, I wonder if flipping that paradigm and asking where we'll end up might reveal something about where we are now.

So: how might VIS (as a venue or field, have your pick) meet its end?

- The people in the community die. The most literal way for this community to end will be for the people who make up this community to die. Perhaps a flood at this year's VIS will wash us all away.
- The community dies. More figuratively, the community (rather than the people in it) might die, such as through infighting about revisions or gate-keeping, or something external might kill the spirit that keeps us together, and we all go separate ways. Perhaps global situations change, and our exclusionary response to them causes people to stop caring about submitting? Perhaps the conference keeps getting hit by hurricanes until people decide that going to a yearly online event isn't worth the trouble. Perhaps a shadowy body somewhere in the upper echelons of our administration will decide none of this is worth it anymore and cancel the whole thing.
- The community gets swallowed, by a whale. There are many whales in the academic sea these days, most prominent among them ML and HCI, but others exist as well. If all of our work is focused within a particular area, we might as well just merge with them.
- By oblivion. Our work may no longer matter, if winds and tides change, we may lose funding and can't pay for this work anymore. Sure, charts are ubiquitous, but so are jokey articles in more serious venues, and they don't sustain 3ish international conferences per year.
- Of boredom. It's been 20 years since Bill Lorenson [8] rang the bell for our death, and we have done little to save ourselves. We have mostly not found new grand challenges to win, we have made poor alliances with other fields. For instance, a VIS23 panel [1] observed that the modal number of TVCG citations in the proceedings of leading ML venues is 0. And finally we have not really embraced our customers more than we were before. Yes, there have been some great tools, but is that enough? When will we stop just retreading old ground with (perhaps Large) new shoes?

In 25 more years, in 2050, where will we be? Will visualization as a research field exist? In the last 25 years, VAST rose and (arguably) fell. As we continue to work on the numerous different avenues we as a field tend to explore, consider: where does this work fit in the life cycle of this field? Are the topics and approaches we have now enough to sustain us?

CLOSING

With that we close the first edition of Charts and Measures, which was to an unambiguous and rousing success. This eclectic collection of articles breaks new ground, offering fresh perspectives untethered from the pedestrian and restricted views that are typically allowed to be on offer at alt.vis. Criticisms were posed, forbidden issues emphasized, and complaints made. These works are truly peerless in their pursuit of pure alterity.

However, this leaves us with a conundrum. How will subsequent editions of this venue maintain unremitting adherence to alternative virtue? Any additional attempts would be necessarily derivative. Is this venue merely to become another alt.vis gradually descending into mainstream acceptance? Is it to become another alt.chi—eventually canceled after forming a vibrant community in the interest of making its host venue cheaper?

Thus we arrive at the only reasonable solution: to maintain the ultimate purity achieved in these pages, we must now close Charts and Measures forever. There will be no more future editions. That is, until someone decides that they need another CV line to support their academic advancement. ♦

ACKNOWLEDGMENTS

We acknowledge Andrew McNutt for his hard work assembling this issue, authoring, reviewing, and acting as expert source for the presented articles. It was supported by grant #1 from Andrew McNutt.

1. A version of this paper that explains the various bits, references, and jokes is available at https://www.mcnutt.in/assets/charts-and-measures-explained.pdf

References

[1] D. Cashman, J. Wang, Q. Wang, D. H. Chau, M. El-Assady, L. Gou, R. Maciejewski, and D. Moritz. How should vis4ml redefine itself in the rapid evolution of ai?, 2023. IEEE VIS Panel.

[2] D3. D3 gallery. https://observablehq.com/@d3/gallery, 2025.

[3] N. Elmqvist. This is not vis: On the perils of gatekeeping in peer review. https://niklaselmqvist.medium.com/this-is-not-vis-3bb50b522c99, 2025.

[4] S. M. Farris. The rise to dominance of genetic model organisms and the decline of curiosity-driven organismal research. PLoS One, 2020.

[5] S. Haroz. Why i don't write a "related work" section. http://steveharoz.com/blog/2015/related-work-section/, 2015.

[6] J. Hathaway. How linkedin's spammy subject line became a new yorker punch line. New York Magazine, 2015.

[7] N. Jardine, B. D. Ondov, N. Elmqvist, and S. Franconeri. The perceptual proxies of visual comparison. IEEE TVCG, 2019.

[8] B. Lorensen. On the death of visualization. In Workshop Visualization Research Challenges, 2004.

[9] A. M. McNutt. No grammar to rule them all: A survey of json-style dsls for visualization. IEEE TVCG, 2022.

[10] A. M. McNutt and R. Chugh. Integrated visualization editing via parameterized declarative templates. In ACM SIGCHI, 2021.

[11] A. M. McNutt, L. Huang, and K. Koenig. Visualization for villainy. alt.vis, 2021.

[12] L. J. Rosenblum. On the birth and impact of the ieee visualization conference. IEEE Computer Graphics and Applications, 45(3):82–94, 2025.

[13] Wikipedia. Model organism, 2025.

[14] K. Wongsuphasawat, Z. Qu, D. Moritz, R. Chang, F. Ouk, A. Anand, J. Mackinlay, B. Howe, and J. Heer. Voyager 2: Augmenting visual analysis with partial view specifications. In ACM SIGCHI, 2017.